

Phase Coexistence for the Hard-Core Model on \mathbb{Z}^2

Antonio Blanca*, Yuxuan Chen[†],
David Galvin[‡], Dana Randall[§] and Prasad Tetali[¶]

November 4, 2016

Abstract

The hard-core model has attracted much attention across several disciplines, representing lattice gases in statistical physics and independent sets in discrete mathematics and computer science. On finite graphs, we are given a parameter λ , and an independent set I arises with probability proportional to $\lambda^{|I|}$. On infinite graphs a Gibbs measure is defined as a suitable limit with the correct conditional probabilities, and we are interested in determining when this limit is unique and when there is phase coexistence, i.e., existence of multiple Gibbs measures.

It has long been conjectured that on \mathbb{Z}^2 this model has a critical value $\lambda_c \approx 3.796$ with the property that if $\lambda < \lambda_c$ then it exhibits uniqueness of phase, while if $\lambda > \lambda_c$ then there is phase coexistence. Much of the work to date on this problem has focused on the regime of uniqueness, with the state of the art being recent work of Sinclair, Srivastava, Štefankovič and Yin showing that there is a unique Gibbs measure for all $\lambda < 2.538$. Here we give the first non-trivial result in the other direction, showing that there are multiple Gibbs measures for all $\lambda > 5.3506$. There is some potential for lowering this bound, but with the methods we are using we cannot hope to replace 5.3506 with anything below about 4.8771.

Our proof begins along the lines of the standard Peierls argument, but we add two innovations. First, following ideas of Kotecký and Randall, we construct an event that distinguishes two boundary conditions and always has long contours associated with it, obviating the need to accurately enumerate short contours. Second, we obtain improved bounds on the number of contours by relating them to a new class of self-avoiding walks on an oriented version of \mathbb{Z}^2 .

*School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332; ablanca@cc.gatech.edu. Research supported in part by the NSF grants CCF-1420934, CCF-1555579 and CCF-1617306.

[†]Computer Science Department, Columbia University, New York NY 10027; yuxuan.chen@columbia.edu. Research supported in part by the Simons Foundation grant 360240.

[‡]Department of Mathematics, University of Notre Dame, Notre Dame IN 46656; dgalvin1@nd.edu. Research supported in part by the Simons Foundation grant 360240 and by the National Security Agency grant NSA H98230-13-1-0248.

[§]School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332; randall@cc.gatech.edu. Research supported in part by the NSF grant CCF-1526900.

[¶]School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332; tetali@math.gatech.edu. Research supported in part by the NSF grant DMS-1407657.

1 Introduction

For a graph G let $\mathcal{I}(G)$ denote the set of independent sets of G . For finite G the *hard-core measure on G with parameter λ* is the measure $\mu_{G,\lambda}$ supported on $\mathcal{I}(G)$ given by $\mu_{G,\lambda}(I) \propto \lambda^{|I|}$ for each $I \in \mathcal{I}(G)$, or equivalently

$$\mu_{G,\lambda}(I) = \frac{\lambda^{|I|}}{\sum_{J \in \mathcal{I}(G)} \lambda^{|J|}}.$$

The hard-core measure is a simple mathematical model of a gas with particles of non-negligible size. The vertices of G are regarded as positions, each of which can be occupied by a particle, subject to the rule that two neighboring sites cannot both be occupied (particles cannot overlap).

On infinite graphs, which may admit infinitely many independent sets, we make sense of the notion of choosing an independent set I with probability proportional to $\lambda^{|I|}$ using the machinery of Gibbs measures. Roughly speaking, these are measures supported on the set of independent sets of G whose conditional restrictions to finite subgraphs agree with the (suitably conditioned) finite hard-core measure.

Formally, let $G = (V, E)$ be infinite and locally finite. For finite $W \subseteq V$, let $\partial_{\text{ext}} W$ be the set of vertices outside W that are adjacent to something in W . Denote by W^+ the subgraph of G induced by $W \cup \partial_{\text{ext}} W$. Let μ be a measure on the set of independent sets of G , equipped with the cylinder σ -algebra (generated by the cylinder events, that is, subsets of the set of independent sets that agree at a particular vertex).

Definition 1.1 *The measure μ is a Gibbs measure with parameter λ for the hard-core model on G if the following condition holds: for all finite $W \subseteq V$ and μ -almost-all independent sets I , the probability that I' agrees with I on W given that it agrees with I off W , with I' drawn according to μ , is the same as the probability that I' agrees with I on W given that it agrees with I on $\partial_{\text{ext}} W$, with I' drawn according to $\mu_{W^+, \lambda}$.*

See, for example, [15] for a very thorough treatment of this topic.

General compactness arguments show that an infinite, locally finite graph G admits at least one Gibbs measure. A central concern of statistical physics (again see [15] for a thorough discussion) is understanding when a particular system — the independent set model in the present setting — exhibits *phase coexistence* (also known as *phase transition*) on a given infinite G , meaning that it admits more than one Gibbs measure.

The canonical (and by far most studied, and physically relevant) case of the hard core measure is that of the usual nearest neighbor graph on the integer grid \mathbb{Z}^d . In this note we specifically consider the two-dimensional grid \mathbb{Z}^2 . Formally this is the graph whose vertex set is the set of pairs $(x, y) \in \mathbb{Z} \times \mathbb{Z}$, with two pairs adjacent if they differ on exactly one coordinate, and differ by ± 1 on that coordinate.

For the classical Ising model, seminal work of Onsager [20] established the precise value ($\beta_c(\mathbb{Z}^2) = \log(1 + \sqrt{2})$) of the critical inverse temperature below which that model exhibits uniqueness of phase and above which it exhibits phase coexistence. Only recently have the analogous values for the (more general) q -state Potts model been established, in work of Beffara and Duminil-Copin [4], settling a more than half-a-century old open problem.

Such precise results for the hard-core model seem far out of reach with currently available methods. It has long been conjectured, though — with computational support, see e.g. [3] — that there is a $\lambda_c \approx 3.796$ such that the hard-core model on \mathbb{Z}^2 exhibits phase coexistence for all $\lambda > \lambda_c$, but not for any $\lambda < \lambda_c$.

Starting with Dobrushin [11] in 1968, both physicists and mathematicians have been developing techniques to approach this conjecture. Most of the attention has focused on establishing ever larger values of λ below which there is uniqueness of phase. The problem has proved to be a fruitful one for the blending of ideas from physics, discrete probability and theoretical computer science, with improvements to our understanding having been made successively by Radulescu and Styer [23], van den Berg and Steif [7], Weitz [29], Restrepo, Shin, Tetali, Vigoda and Yang [25], and Vera, Vigoda and Yang [28], among others. The state of the art is recent work of Sinclair, Srivastava, Štefankovič and Yin [26], building on the novel ideas of Weitz, which establishes that there is a unique Gibbs measure for all $\lambda < 2.538$.

Much less is known about the regime of phase coexistence. Dobrushin [11] established that there is a $C > 0$ such that for all $\lambda > C$ there are multiple Gibbs measures. He did not explicitly calculate C , but around the time of the writing of [9] Borgs [8] reported that a direct implementation of Dobrushin’s argument would yield a value of C strictly greater than 80. A later computation by the third author suggested that the actual consequence of the argument was more like $C \approx 300$.

Our main aim in this paper is to give a reasonable upper bound on λ_c (if it exists).

Theorem 1.1 *The hard-core model on \mathbb{Z}^2 with activity λ admits multiple Gibbs measures for all $\lambda > 5.3506$.*

An intuition for the meaning of multiple Gibbs measures can be gleaned from the following recipe for producing them. For an independent set I and a finite set $W \subseteq V$, let $\mathcal{I}^I(W)$ be the (also finite) set of independent sets that agree with I off W . Fix I and a nested sequence $(W_i)_{i=1}^\infty$ of finite subsets of V satisfying $\cup_i W_i = V$. For each i let μ_i^I be the measure on $\mathcal{I}^I(W_i)$ in which each J is selected with probability proportional to $\lambda^{|J \cap W_i|}$. Any (weak) subsequential limit of the μ_i^I ’s (and by compactness there must be at least one such) is a Gibbs measure. This fact was originally proved, in a much more general context, by Dobrushin [12]; see for example [7] for a treatment specific to the hard-core model on the lattice, or [10, Theorem 3.5] for a simple proof in the slightly more general context of graph homomorphism models.

From this recipe we see that an interpretation of the existence of multiple Gibbs measures on \mathbb{Z}^2 is that the local behavior of a randomly chosen independent set in a box can be made to depend on a boundary condition imposed on the box, even in the limit as the size of the box grows to infinity. This leads to what turns out to be the standard approach to showing multiple Gibbs measures, which is to consider the limiting measures corresponding to two different boundary conditions on boxes in the lattice centered at the origin, and to find a statistic that separates these two limits. For the hard-core model, it suffices (see [7]) to compare the even boundary condition — all vertices on the boundary of a box at an even distance from the origin are occupied — and its counterpart odd boundary condition, and the distinguishing statistic is typically the occupation of the origin. Under odd boundary condition the origin should be unlikely to be occupied, since independent sets with odd boundary and (even) origin occupied must have a contour — a two-layer thick unoccupied ring of vertices separating an inner region around the origin that is in “even phase” from

an outer region near the boundary that is in “odd phase”. For large enough λ , such an unoccupied layer is costly, and so such configurations are unlikely. This is essentially the Peierls argument for phase coexistence, and was the approach taken by Dobrushin [11].

As we will see presently, the effectiveness of the Peierls argument is driven by the number of contours of each possible length – better upper bounds on the number of contours translate directly to better upper bounds on λ_c . Previous work of other researchers on phase coexistence in the hard-core model on \mathbb{Z}^2 had viewed contours as simple polygons in \mathbb{Z}^2 , which are closely related to the very well studied family of self-avoiding walks. While this is essentially the best possible point of view when applying the Peierls argument on the Ising model, it is far from optimal for the hard-core model. One of the two contributions of the present paper is the realization that hard-core contours, if appropriately defined, can be viewed as simple polygons in the oriented Manhattan lattice, in which edges of \mathbb{Z}^2 that are parallel to the x -axis (respectively, the y -axis) are oriented positively if their y -coordinate (respectively, x -coordinate) is even, and negatively otherwise, with the additional constraint that contours cannot make two consecutive turns. The number of such polygons can be understood by analyzing a new class of self-avoiding walks, that we refer to as *taxi walks*. The number of taxi walks turns out to be significantly smaller than the number of ordinary self-avoiding walks, leading to significantly better bounds on λ_c .

There is a single number $\mu_{\text{taxi}} > 0$, the taxi walk connective constant, that asymptotically controls the number c_n of taxi walks of length n , in the sense that $c_n = \mu_{\text{taxi}}^n f_{\text{taxi}}(n)$ with $f_{\text{taxi}}(n)$ sub-exponential. Adapting methods of Goulden and Jackson [16] we obtain estimates on μ_{taxi} , giving us good understanding of c_n for large n . The sub-exponential correction makes it difficult to control c_n for small n , however, presenting a major stumbling block to the effectiveness of the Peierls argument as we have just described it. Using the statistic “occupation of origin” to distinguish the two boundary conditions, one inevitably has to control c_n for both small and large n . The lack of precise information about the number of short contours leads to discrepancies between lower limits of arguments and actual bounds, such as that between $C = 80$ (lower bound on C) and $C \approx 300$ from a direct implementation of Dobrushin’s argument for phase coexistence on \mathbb{Z}^2 discussed earlier.

The second contribution of the present paper is the idea of using an event to distinguish the two boundary conditions that has the property that every independent set in the event has associated with it a long contour. This allows us to focus exclusively on the asymptotic growth rate of contours/taxi walks, and obviates the need for an analysis of short contours. The immediate result of this breakthrough is that the actual consequences of our arguments agree exactly with their natural lower limits.

In an earlier version of this work [5] we used the technology of fault lines, introduced in [24] to define a distinguishing event. After a talk by one of the authors, Kotecký pointed out an alternate approach [18]. Consider a box B of fixed size centered at the origin, and say that an independent set is even (respectively, odd) on the box if every vertex in B at an even (respectively, odd) distance from the origin is either in the independent set, or potentially could be in the sense that none of its neighbors are. Then consider the event that a randomly chosen independent set drawn from a much larger box (with boundary condition) is even on B , conditioned on the event that it is either even or odd on B . Running the Peierls argument on this event leads to contours that completely encircle B , and so can be made arbitrarily long by choosing B to be sufficiently large.

In Section 2 we introduce the notion of taxi walks and the taxi walk connective constant, that will be key to the precise theorem we prove (Theorem 2.3), from which Theorem 1.1 follows via some numerical computation. The proof of Theorem 2.3 is then given in Section 3. In Section 4 we give the details of our upper and lower bounds on the taxi walk connective constant, and we conclude in Section 5 with some remarks.

2 Taxi walks

Let $\vec{\mathbb{Z}}^2$ be an orientation of \mathbb{Z}^2 in which an edge parallel to the x -axis (respectively, y -axis) is oriented in the positive x -direction if its y -coordinate is even (respectively, oriented in the positive y -direction if its x -coordinate is even), and is oriented in the negative direction otherwise. It is common to refer to $\vec{\mathbb{Z}}^2$ as the *Manhattan lattice*: streets are horizontal, with even numbered streets oriented west to east and odd numbered streets oriented east to west, and avenues are vertical, with even numbered avenues oriented south to north and odd numbered avenues oriented north to south.

Definition 2.1 A self-avoiding walk in $\vec{\mathbb{Z}}^2$ of length n starting at vertex v_0 is a sequence v_0, v_1, \dots, v_n of distinct vertices with, for each $i = 1, \dots, n$, $v_{i-1}v_i$ an edge of $\vec{\mathbb{Z}}^2$ oriented from v_{i-1} to v_i . The walk turns at v_i ($1 \leq i \leq n-1$) if edges $v_{i-1}v_i$ and v_iv_{i+1} are perpendicular, and goes straight if these edges are parallel. A taxi walk is a self-avoiding walk in $\vec{\mathbb{Z}}^2$ that does not turn at two consecutive vertices.

We call these taxi walks because a savvy passenger in a Manhattan cab would be suspicious if the cab took two consecutive turns.

Let c_n be the number of taxi walks of length n starting at the origin. A critical step in our arguments will be bounding c_n . An easy upper bound is $c_n \leq 2^n$, since there are always at most two ways to extend a taxi walk of length $n-1$, and an easy lower bound is $2^{n/2} \leq c_n$ (for even n) and $2^{(n+1)/2} \leq c_n$ (for odd n), since walks that always take two steps north at a time or two steps east at a time can always be extended in exactly two ways. With very little extra work we can get a significantly better upper bound.

Lemma 2.1 $c_n = O(((1 + \sqrt{5})/2)^n)$.

Proof: A taxi walk v_0, v_1, \dots, v_n (with v_0 the origin) can be encoded by a pair (a, σ) , where $a \in \{N, E\}$ and σ is a sequence of length $n-1$ over alphabet $\{s, t\}$, as follows: if $v_1 = (0, 1)$ then $a = N$ and if $v_1 = (1, 0)$ then $a = E$, and if the walk goes straight at v_i then the i th entry of σ is s , whereas it is t if the walk turns at v_i . Distinct taxi walks evidently get distinct codes. It is well known that the number of sequences of length $n-1$ over alphabet $\{s, t\}$ without consecutive occurrences of the character t is the $(n+1)$ st Fibonacci number f_{n+1} (defined by $f_0 = 0$, $f_1 = 1$, $f_n = f_{n-1} + f_{n-2}$ for $n \geq 2$). It follows that $c_n \leq 2f_{n+1} = O(((1 + \sqrt{5})/2)^n)$. \square

Using more sophisticated tools we can improve our bounds. In what follows we say that a function $f(n)$ defined on positive integers grows *sub-exponentially* if for all $\varepsilon > 0$ there is $n(\varepsilon)$ such that for all $n > n(\varepsilon)$ we have $f(n) < (1 + \varepsilon)^n$.

Theorem 2.2 *There is a constant μ_{taxi} (the taxi walk connective constant) with $1.55701 < \mu_{\text{taxi}} < 1.58746$ and $4.8771 < \mu_{\text{taxi}}^4 - 1 < 5.3506$, and a function $f_{\text{taxi}}(n)$ that grows sub-exponentially, such that $c_n = f_{\text{taxi}}(n)\mu_{\text{taxi}}^n$ for all n .*

We defer the proof of Theorem 2.2 to Section 4. We end this section with a precise statement of our main theorem, from which Theorem 1.1 follows via Theorem 2.2.

Theorem 2.3 *The hard-core model on \mathbb{Z}^2 with activity λ admits multiple Gibbs measures for all $\lambda > \mu_{\text{taxi}}^4 - 1$.*

Our bounds on $\mu_{\text{taxi}}^4 - 1$ involve the theory of irreducible bridges and the Goulden-Jackson cluster method, as well as extensive computation, but the proof of the existence of μ_{taxi} and $f_{\text{taxi}}(n)$ at the beginning of Section 4 and the bounds $\sqrt{2} \leq \mu_{\text{taxi}} \leq (1 + \sqrt{5})/2$ in this section are straightforward. A consequence of this is that if we wish to avoid using irreducible bridges, the Goulden-Jackson cluster method and computer-aided computations, we have via Theorem 2.3 and Lemma 2.1 a weaker version of Theorem 1.1 that is still significantly better than any previous result; namely, that the hard-core model on \mathbb{Z}^2 with activity λ admits multiple Gibbs measures for all $\lambda > (5 + 3\sqrt{5})/2$, and so for all $\lambda > 5.8542$.

A consequence of our lower bound on μ_{taxi} is that our present approach to phase co-existence cannot prove anything better than $\lambda_c \leq 4.8771$. A computation by Pantone [21] using the method of differential approximants [22] on the sequence $(c_n)_{n=1}^{60}$ suggests $\mu_{\text{taxi}} \in [1.57376, 1.57378]$, so that the limit of the present approach may in fact be more like 5.134. Note that a similar situation exists for lower bounds on λ_c : Vera, Vigoda and Yang [28] showed $\lambda_c \geq 2.48$, but they also observe that the methods they use cannot prove anything better than $\lambda_c \geq 3.4$, which falls below the conjectured value of λ_c , around 3.796 (established by Baxter, Enting and Tsang [3]); and although Sinclair, Srivastava, Štefankovič and Yin do not rule out the possibility that the approach described in [26] might eventually bring the lower bound all the way up to the critical value, they express some skepticism.

3 Proof of phase coexistence (Theorem 2.3)

Let $\lambda > \mu_{\text{taxi}}^4 - 1$ be fixed. Our argument will depend on a parameter $m = m(\lambda)$ whose value will be specified later.

Let \mathcal{E} denote the set of *even* vertices of \mathbb{Z}^2 — those vertices (x, y) with $x + y$ even — and let \mathcal{O} denote the complementary set of odd vertices; note that these are both independent sets. Let U_n be the box $\{-n, -(n-1), \dots, n-1, n\}^2$. When I, J are independent sets such that $I \subseteq J$, we say that J *extends* I . Let \mathcal{J}_n^e be the set of independent sets that extend $\mathcal{E} \setminus U_n$, and let μ_n^e be the probability distribution supported on \mathcal{J}_n^e in which each set I is selected with probability proportional to $\lambda^{|I \cap U_n|}$. Define μ_n^o analogously (with “even” everywhere replaced by “odd”).

Say that an independent set I in \mathbb{Z}^2 is *m-even* if for every even $x \in U_m$, none of the four neighbors of x is in I , and define *m-odd* analogously. Say that I is *m-homogeneous* if it is either *m-odd* or *m-even*.

Now fix $n > m$ and let \mathcal{E}_m be the event that an independent set is *m-even*, \mathcal{O}_m the event that it is *m-odd*, and \mathcal{H}_m the event that it is *m-homogeneous*. Note that all of these

events are in the cylinder σ -algebra. We will establish the following conditional probability inequality for all $n \geq m$ and m large enough:

$$\mu_n^e(\mathcal{O}_m | \mathcal{H}_m) < 1/3. \quad (1)$$

Reversing the roles of odd and even throughout the proof, we will also get

$$\mu_n^o(\mathcal{E}_m | \mathcal{H}_m) < 1/3$$

and so

$$\mu_n^o(\mathcal{O}_m | \mathcal{H}_m) > 2/3.$$

It follows that if μ^e is any Gibbs measure obtained as a weak subsequential limit of the μ_n^e 's, and μ^o is any obtained from the μ_n^o 's, then

$$\mu^e(\mathcal{O}_m | \mathcal{H}_m) \leq 1/3 < 2/3 \leq \mu^o(\mathcal{O}_m | \mathcal{H}_m);$$

consequently, μ^e and μ^o are distinct Gibbs measures.

Write \mathcal{B}_n^e for the set of independent sets in \mathbb{Z}^2 that extend $\mathcal{E} \setminus U_n$ and are m -odd, and write \mathcal{A}_n^e for the set of independent sets in \mathbb{Z}^2 that extend $\mathcal{E} \setminus U_n$ and are m -homogeneous (so $\mathcal{B}_n^e \subseteq \mathcal{A}_n^e \subseteq \mathcal{J}_n^e$).

For $I \in \mathcal{A}_n^e$ set $w_\lambda(I) = \lambda^{|I \cap U_n|}$, and for a set \mathcal{C} of independent sets in \mathcal{A}_n^e let $w_\lambda(\mathcal{C})$ denotes $\sum_{I \in \mathcal{C}} w_\lambda(I)$. To establish (1) it is enough to show

$$\frac{\mu^e(\mathcal{O}_m)}{\mu^e(\mathcal{H}_m)} = \frac{w_\lambda(\mathcal{B}_n^e)}{w_\lambda(\mathcal{A}_n^e)} < 1/3. \quad (2)$$

The intuition here is that if I is conditioned to agree with \mathcal{E} outside U_n , then under the extra condition that it is m -homogeneous it is far less likely to be m -odd than m -even.

We will establish (2) by constructing, for each $I \in \mathcal{B}_n^e$, a collection $\varphi(I) \subseteq \mathcal{A}_n^e$, together with a flow function $f(I, J)$ supported on $\{(I, J) : I \in \mathcal{B}_n^e, J \in \varphi(I)\}$ that satisfies

$$\sum_{J \in \varphi(I)} f(I, J) = 1 \quad \text{for each } I \in \mathcal{B}_n^e \quad (3)$$

and

$$\sum_{I: J \in \varphi(I)} \lambda^{|I \cap U_n| - |J \cap U_n|} f(I, J) < 1/3 \quad \text{for each } J \in \mathcal{A}_n^e. \quad (4)$$

This gives (2), via

$$\begin{aligned} w_\lambda(\mathcal{B}_n^e) &= \sum_{I \in \mathcal{B}_n^e} \lambda^{|I \cap U_n|} \\ &= \sum_{I \in \mathcal{B}_n^e} \sum_{J \in \varphi(I)} \lambda^{|I \cap U_n|} f(I, J) \\ &= \sum_{J \in \mathcal{A}_n^e} \lambda^{|J \cap U_n|} \sum_{I: J \in \varphi(I)} \lambda^{|I \cap U_n| - |J \cap U_n|} f(I, J) \\ &< w_\lambda(\mathcal{A}_n^e)/3. \end{aligned}$$

To construct φ we will use the fact that $I \in \mathcal{B}_n^e$ is in *even phase* (predominantly even-occupied) outside U_n , but because I is m -odd, it is not in even phase close to U_m ; so there must be a *contour* — an unoccupied ring of vertices — marking the extent of the even phase inside U_n .

We will proceed in two stages. In Section 3.1, we explain how such a contour can be explicitly constructed, and establish the various properties of the construction that we will need. (To aid readability we defer proofs of many of these properties to Sections 3.3 and 3.4.) In Section 3.2 we describe and analyze the standard Peierls argument, which involves modifying I inside the contour to create $\varphi(I)$ satisfying (3) and (4) (for suitable choice of the flow function f), showing that being m -odd is unlikely, conditioned on being m -homogeneous, under even boundary condition.

3.1 The contour and its properties

Fix $I \in \mathcal{B}_n^e$ (so I includes all even vertices outside U_n , and is m -odd). Let I' consist of I together with each odd vertex that has none of its neighbors in I ; notice that $I' \in \mathcal{B}_n^e$, that it is completely determined by I , and that it includes all odd vertices of U_m (since I is m -odd). The point of passing from I to I' is that in doing so, we ensure that the contour we construct fully encircles U_m and hence has length at least on the order of m . We now describe how to associate with I' a set $\gamma(I)$ of edges of \mathbb{Z}^2 , which we will refer to as the *contour* associated with I . The same construction was used in [9] and [13], and a very similar construction appeared in [14].

We begin with a brief reminder of some graph theory notation. Given $S \subseteq \mathbb{Z}^2$ the subgraph of \mathbb{Z}^2 induced by S is the graph with vertex set S in which two vertices are adjacent if and only if there are adjacent in \mathbb{Z}^2 . We will abuse notation somewhat and refer to this graph simply as S . Given distinct vertices $u, v \in S$ a *path* in S from u to v (a u - v path) is a sequence $u = u_0, u_1, \dots, u_k = v$ of distinct vertices of S with u_i and u_{i+1} adjacent for each $i = 0, \dots, k-1$. Given a vertex $u \in S$ and a subset $T \subseteq S$ of vertices with $u \notin T$ a path in S from u to T (a u - T path) is any path from u to v with $v \in T$. A *component* of S is a subset $C \subseteq S$ of vertices with the property that for any pair of distinct vertices $u, v \in C$ there is a path in S from u to v , and that is maximal with respect to this property.

Let $(I^\circ)^+$ be the set of odd vertices in I' together with their neighbors. By the m -oddness of I , U_m is contained in a single component of the graph induced by $(I^\circ)^+$; let R be that component. Note that because I' extends $\mathcal{E} \setminus U_n$, R is finite, and specifically $R \subseteq U_n$.

The following property, which says that in leaving R one always goes from an unoccupied even vertex to an unoccupied odd vertex (in fact, an odd vertex outside I'), is evident from the construction of R .

$$\text{If } uv \text{ is an edge with } u \in R \text{ and } v \notin R \text{ then } u \in \mathcal{E} \setminus I \text{ and } v \in \mathcal{O} \setminus I'. \quad (5)$$

Define $\gamma = \gamma(I)$ to be the set of edges uv with $u \in R$ and $v \notin R$, and such that there is a path in \mathbb{Z}^2 from v to $\mathbb{Z}^2 \setminus U_n$ that avoids R . (This last condition has the effect of removing any holes that R may have.)

We now introduce a graph \mathbb{Z}_\diamond^2 that may be thought of dual to \mathbb{Z}^2 . The vertices of \mathbb{Z}_\diamond^2 are the midpoints of edges in \mathbb{Z}^2 , and two such vertices are adjacent if the associated edges in \mathbb{Z}^2 are incident and perpendicular. Note that \mathbb{Z}_\diamond^2 is a rotated, dilated, translated copy of \mathbb{Z}^2 .

Each edge in γ is a vertex in \mathbb{Z}_\diamond^2 , and so we may specify γ by specifying a subgraph of \mathbb{Z}_\diamond^2 whose vertex set is γ . In what follows we describe a way to specify one such subgraph, which will turn out to be a cycle (a connected 2-regular graph) in \mathbb{Z}_\diamond^2 that separates R from $\mathbb{Z}^2 \setminus U_n$, and moreover has a particular structure related to taxi walks.

To construct the subgraph, which we will call Γ and also sometimes refer to as the contour associated with I , consider an arbitrary $uv \in \gamma$ with $u \in R$ and $v \notin R$ (so $u \in \mathcal{E}$ and $v \in \mathcal{O}$). The edge uv forms a side of two 1-by-1 squares in \mathbb{Z}^2 . Let $uvst$ be one such, with $s \in \mathcal{E}$ and $t \in \mathcal{O}$ (so each of uv , vs , st and tu are edges of \mathbb{Z}^2). Viewed as vertices of \mathbb{Z}_\diamond^2 , uv has two neighbors among the edges uv , vs , st and tu , namely tu and vs . If $s \in R$, then the directed edge $uv \rightarrow vs$ is added to Γ , and if $s \notin R$ then the directed edge $uv \rightarrow tu$ is added to Γ . We proceed in the same way with the other 1-by-1 square in \mathbb{Z}^2 that uv forms one side of. (Thus Γ is initially a directed graph; presently we will modify it slightly to create an undirected graph.)

It will be helpful for subsequent arguments to view the construction above via the following case-by-case analysis, which is easily seen to be equivalent. There are four cases, depending on the statuses of s and t with regards to membership of R .

Case i, $t \in R$ and $s \notin R$: This case cannot occur since it gives rise to the edge st with $s \in R$ and $t \notin R$ but $s \in \mathcal{E}$, violating (5).

Case ii, $t \notin R$ and $s \notin R$: In this case, of tu and vs only tu is in γ (to see that $tu \in \gamma$ note that there is a path in \mathbb{Z}^2 from t to $\mathbb{Z}^2 \setminus U_n$ that avoids R , that starts tsv and then continues along any R -avoiding path from v to $\mathbb{Z}^2 \setminus U_n$), and we put the directed edge (in \mathbb{Z}_\diamond^2) $uv \rightarrow tu$ in Γ .

Case iii, $t \in R$ and $s \in R$: In this case, of tu and vs only vs is in γ , and we put the edge $uv \rightarrow vs$ in Γ .

Case iv, $t \notin R$ and $s \in R$: In this case, vs is evidently in γ , and ut may or not be (depending on whether there is an R -avoiding path in \mathbb{Z}^2 from t to $\mathbb{Z}^2 \setminus U_n$), and we put the edge $uv \rightarrow vs$ in Γ (but *not* the edge from uv to tu , even if $tu \in \gamma$).

What we have constructed so far is a *directed* graph on the set of vertices in \mathbb{Z}_\diamond^2 corresponding to edges in γ , with the property that every vertex has out-degree two. But notice that the construction of directed edges is symmetric: if we put the edge $uv \rightarrow tu$ (say) in Γ then the construction also mandates putting in the edge $tu \rightarrow uv$. So underlying the 2-out-regular directed graph is a 2-regular undirected graph (a union of cycles) and it is this we take as Γ . Note that Γ determines γ , since its vertex set is exactly γ ; but Γ is not (necessarily) the subgraph of \mathbb{Z}_\diamond^2 induced by γ , because not all edges from \mathbb{Z}_\diamond^2 with endpoints in γ are in Γ (see Case iv above). Note also that because Γ is 2-regular, the size of both its vertex set and its edge set is $|\gamma|$.

If we draw Γ in \mathbb{Z}^2 , using straight-line segments joining midpoints of edges of γ to represent edges of Γ , then each component of Γ is a simple closed \mathbb{Z}^2 -avoiding polygon in \mathbb{R}^2 , and so encloses a finite interior (with an infinite exterior). We refer to the vertices of \mathbb{Z}^2 that are in the interior as the *vertex interior* of the component, and to all other vertices of \mathbb{Z}^2 as the *vertex exterior*. A basic fact is the following.

Lemma 3.1 *If $uv \in \gamma$ with $u \in R$ and $v \notin R$, then u is in the vertex interior of the component of uv in Γ , and v is in the vertex exterior.*

Proof: Assume without loss of generality that uv is parallel to the y -axis. By construction

Γ has an edge to the right of uv that ends at the midpoint of uv , and one to the left that starts at that point, so the edge uv crosses Γ . This establishes that one of u, v is in the vertex interior of the component of uv in Γ , and the other is in the vertex exterior. If v is in the vertex interior then any $v\text{-}\mathbb{Z}^2 \setminus U_n$ path in \mathbb{Z}^2 must cross Γ and so meet R , contradicting the construction of γ . So v is in the vertex exterior and u is in the vertex interior. \square

So far we have established that Γ is a union of cycles; Lemma 3.1 is a key ingredient in proving our first important fact about Γ ; namely, that it has a single component.

Lemma 3.2 *The graph Γ is a cycle.*

Proof: Assume, for a contradiction, that Γ has distinct components C_1 and C_2 . It cannot be the case that one of these, C_1 say, encloses the other. For if $uv \in \gamma$ crosses C_2 , with $v \notin R$, then there must (by the definition of γ) be a $v\text{-}\mathbb{Z}^2 \setminus U_n$ path in \mathbb{Z}^2 that avoids R ; but any such path must cross C_1 , and so meet R , a contradiction.

So it must be the case that C_1 and C_2 have disjoint vertex interiors. Let uv cross C_1 and $u'v'$ cross C_2 , with u and u' the interior vertices. Via Lemma 3.1 this gives an immediate contradiction: every $u\text{-}u'$ path in \mathbb{Z}^2 must use an edge of γ and so leave R , contradicting the connectivity of R . \square

The following facts about the structure of Γ (Lemmas 3.3 and 3.4) will be used to complete the proof of Theorem 2.3. We defer the proof of Lemma 3.3 to Section 3.3, and that of Lemma 3.4, which crucially depends on the connection between contours and taxi walks (essentially, a contour is a closed taxi walk) to Section 3.4.

Lemma 3.3 $|\Gamma| \geq 2\sqrt{2}m$ and is a multiple of 4.

Let \mathcal{C}_ℓ^m be the collection of all Γ with $|\Gamma| = 4\ell$ that arise in the above-described construction, as I runs over \mathcal{B}_n^e , and let $\mathcal{C}^m = \cup_\ell \mathcal{C}_\ell^m$.

Lemma 3.4 *There is a function $g(\ell)$ that grows subexponentially such that $|\mathcal{C}_\ell^m| \leq g(\ell)\mu_{\text{taxi}}^{4\ell}$, where μ_{taxi} is the connective constant of taxi walks.*

The Peierls argument that we will use in Section 3.2 involves the *shift* operation. Essentially this is a shifting, by one lattice unit, of all the vertices of I that are enclosed by Γ , while leaving the remainder of I unchanged. The content of Lemma 3.5 below is that this allows I to be augmented substantially, leading to the a weight-increasing map from \mathcal{B}_n^e to \mathcal{A}_n^e that allows for the construction of a flow function satisfying (3) and (4).

Let W be the vertex interior of Γ and W' the vertex exterior. For $v \in \mathbb{Z}^2$, and $s \in \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$, let $\sigma_s(v) = v + s$. Let $I_s = (I \cap W') \cup \{\sigma_s(v) : v \in I \cap W\}$ — I_s may be thought of as the result of shifting I one unit in the s direction within W , while leaving it unchanged outside W . Let $\tilde{I}_s = \{v \in W : \sigma_s^{-1}(v) \in W'\}$ — we may think of \tilde{I}_s as the set of vertices in W with the property that their preimage under the map σ_s is outside W ; note that $\sigma_s^{-1}(v) = v - s$. Finally, let $I_s'' = I_s \cup \tilde{I}_s$.

Lemma 3.5 *The shifted set I_s is an independent set, with $|I_s| = |I|$. Moreover, the augmented shifted set I_s'' is an independent set, with $|I_s''| = |I_s| + |\tilde{I}_s|$. Finally, there is a choice of s for which $|\tilde{I}_s| \geq |\gamma|/4$.*

The proof of Lemma 3.5 uses standard ideas (see, e.g., [14, Proposition 2.12], [13, Lemma 4.1] and [9, proof of Lemma 6]). The same is true for our next lemma (see, e.g., [14, equation (15)]). For completeness we furnish proofs in Section 3.3.

Lemma 3.6 *If $I \in \mathcal{B}_n^e$ has associated contour Γ , and $J = I_s \cup S$ where $S \subseteq \tilde{I}_s$ and s is one of the four possible shift directions, then I is completely determined by J , s and γ .*

We also need one more lemma, that says that after shifting we go from m -odd independent set to an m -homogeneous set. The proof appears in Section 3.3.

Lemma 3.7 *If $I \in \mathcal{B}_n^e$ then $I_s'' \in \mathcal{A}_n^e$.*

3.2 The Peierls argument for phase coexistence

We are now in a position to define $\varphi(I)$ and $f(I, J)$ for $I \in \mathcal{B}_n^e$ and $J \in \varphi(I)$. First, choose an $s \in \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$ for which $|\tilde{I}_s| \geq |\gamma|/4$ (by Lemma 3.5 there is such an s ; choose, for example, the first such that works in some arbitrary ordering). Next, set

$$\varphi(I) = \{I_s \cup S : S \subseteq \tilde{I}_s\}$$

(by Lemma 3.7 we have $\varphi(I) \subseteq \mathcal{A}_n^e$). Finally, for $I \in \mathcal{B}_n^e$ and $J \in \varphi(I)$ set

$$f(I, J) = \frac{\lambda^{|S|}}{(1 + \lambda)^{|\tilde{I}_s|}}.$$

For this choice of f we have

$$\sum_{J \in \varphi(I)} f(I, J) = \sum_{S \subseteq \tilde{I}_s} \frac{\lambda^{|S|}}{(1 + \lambda)^{|\tilde{I}_s|}} = \frac{1}{(1 + \lambda)^{|\tilde{I}_s|}} \sum_{k=0}^{|\tilde{I}_s|} \binom{|\tilde{I}_s|}{k} \lambda^k = 1,$$

so (3) is established, and it only remains to verify (4).

We now present the Peierls argument that verifies (4) and completes the proof of phase coexistence. Recall that $\lambda > \mu_{\text{taxi}}^4 - 1$ has been given. Choose $\mu > 0$ to be such that $\mu^4 - 1$ is the midpoint of $[\mu_{\text{taxi}}^4 - 1, \lambda]$. Choose m sufficiently large that

$$g(\ell) \mu_{\text{taxi}}^{4\ell} < \mu^{4\ell} \tag{6}$$

for all $\ell \geq \sqrt{2}m/2$, where $g(\ell)$ is the subexponential function from Lemma 3.4 (this inequality holds for all sufficiently large $m = m(\lambda)$ since $\mu > \mu_{\text{taxi}}$, and depends only on λ), and also that

$$\sum_{\ell \geq \sqrt{2}m/2} \frac{\mu^{4\ell}}{(1 + \lambda)^\ell} < 1/3 \tag{7}$$

(this inequality holds for all sufficiently large $m = m(\lambda)$ since $1 + \lambda > \mu^4$).

For $n > m$, fix $J \in \mathcal{A}_n^e$. From the definitions of $\varphi(I)$ and f , we have that

$$\sum_{I: J \in \varphi(I)} \lambda^{|I| - |J|} f(I, J) = \sum_{I: J \in \varphi(I)} \frac{1}{(1 + \lambda)^{|\tilde{I}_s|}} \leq \sum_{I: J \in \varphi(I)} \frac{1}{(1 + \lambda)^{|\gamma(I)|/4}}, \tag{8}$$

since $|\tilde{I}_s| \geq |\gamma(I)|/4$ by Lemma 3.5. Bearing Lemma 3.6 in mind, for each $\Gamma \in \mathcal{C}^m$ there are at most four $I \in \mathcal{B}_n^e$ for which $J \in \varphi(I)$ (at most one for each shift $s \in \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$). Hence,

$$\begin{aligned}
\sum_{I: J \in \varphi(I)} \lambda^{|I|-|J|} f(I, J) &\leq 4 \sum_{\Gamma \in \mathcal{C}^m} \frac{1}{(1 + \lambda)^{|\Gamma|/4}} \\
&= 4 \sum_{4\ell \geq 2\sqrt{2}m} \frac{|\mathcal{C}_\ell^m|}{(1 + \lambda)^\ell} \\
&\leq \sum_{\ell \geq \sqrt{2}m/2} \frac{g(\ell) \mu_{\text{taxi}}^{4\ell}}{(1 + \lambda)^\ell} \\
&< \sum_{\ell \geq \sqrt{2}m/2} \frac{\mu^{4\ell}}{(1 + \lambda)^\ell} \\
&< 1/3,
\end{aligned} \tag{9}$$

where the first inequality uses (8), we use Lemma 3.3 in (9), the second inequality uses Lemma 3.4, and the remaining inequalities follow from our choice of m (specifically using (6) and (7)). The proof is now complete. \square

3.3 Proofs of contours facts

In this section we begin wrapping up the proof of Theorem 2.3 by providing the proofs of Lemmas 3.3, 3.5, 3.6 and 3.7.

We will need a basic fact about Γ that comes immediately from the construction.

Lemma 3.8 *If $\{a, b, c, d\}$ are the vertices of a 1-by-1 square in \mathbb{Z}^2 (with ab, bc, cd and da the edges of \mathbb{Z}^2), then in Γ it is not possible for bc to be adjacent to both ab and cd .*

Proof of Lemma 3.3: Since the interior of Γ contains a vertex with x -coordinate m (along the top of U_m) and one with x -coordinate $-m$ (along the bottom), and each edge of Γ spans a distance of $1/\sqrt{2}$ in the x -direction, it follows that Γ must have at least $2\sqrt{2}m$ edges.

To argue about the length of Γ , we view it as a simple closed \mathbb{Z}^2 -avoiding polygon in \mathbb{R}^2 , in the manner described before the statement of Lemma 3.1, and consider traversing this polygon in a clockwise direction starting from an arbitrarily chosen point P that is the midpoint of the form $(a_x, a_y - 1/2)$ of an edge of γ , with a_x, a_y integers. We traverse in steps of length $\sqrt{2}/2$, which corresponds to moving from the midpoint of one edge of γ to the midpoint of an adjacent (and perpendicular) edge.

A complete traverse of the polygon consists of x_{\searrow} steps oriented southeast (parallel to the edge from $(0, 0)$ to $(-1, -1)$), x_{\nwarrow} steps oriented northwest, x_{\swarrow} steps oriented southwest and x_{\nearrow} steps oriented northeast, and because Γ is closed we have $x_{\searrow} = x_{\nwarrow}$ and $x_{\swarrow} = x_{\nearrow}$.

Starting at P , a point in \mathbb{R}^2 of the form $(a_x, a_y - 1/2)$ with a_x, a_y integers, after two steps we return to a point of this form, having passed through a point of the form $(a'_x - 1/2, a'_y)$ with a'_x, a'_y integers. These two steps must be one of: southwest followed by southeast or vice

versa; northwest followed by northeast or vice versa; or two steps in the same direction. (All other possibilities, such as southwest followed by northwest, are ruled out by Lemma 3.8).

Write $x_{\swarrow\searrow}$ for the total number (over the entire polygon) of pairs of steps of the kind just described that consist of southwest followed by southeast, and write $x_{\swarrow\swarrow}$, $x_{\nwarrow\swarrow}$, $x_{\nwarrow\nwarrow}$, $x_{\searrow\swarrow}$, $x_{\searrow\nwarrow}$, $x_{\swarrow\swarrow}$ and $x_{\nwarrow\nwarrow}$ for the count of the other possible pairs. Using $x_{\swarrow} = x_{\nwarrow}$ we get

$$x_{\swarrow\searrow} + x_{\nwarrow\swarrow} + 2x_{\nwarrow\nwarrow} = x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + 2x_{\nwarrow\nwarrow}$$

and using $x_{\swarrow} = x_{\nwarrow}$ we get

$$x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + 2x_{\nwarrow\nwarrow} = x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + 2x_{\nwarrow\nwarrow}.$$

Combining (and dividing by 2) we get

$$x_{\swarrow\searrow} + x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\nwarrow} = x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\nwarrow}.$$

It follows that $x_{\swarrow\searrow} + x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\swarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\nwarrow} + x_{\nwarrow\nwarrow}$ is even, and so $|\Gamma|$, being twice this sum, is a multiple of 4. \square

That 4 divides $|\Gamma|$ could also be read out of [9, Lemma 5]; we give a self-contained proof above to avoid a lengthy detour matching our notation to that of [9].

Proof of Lemma 3.5: We begin with the final statement. Let γ_s be the set of edges in γ of the form uv with $u \in W$, $v \notin W$ and $v = u - s$. Note that $\gamma = \cup_s \gamma_s$, so there is a choice of s for which $|\gamma_s| \geq |\gamma|/4$. Now the map from γ_s to \tilde{I}_s that sends uv to u is injective (for each $u \in \tilde{I}_s$ there is a unique v such that $uv \in \gamma_s$, namely $u - s$), so $|\tilde{I}_s| \geq |\gamma_s| \geq |\gamma|/4$.

To show $|I_s| = |I|$ consider the map φ from I to I_s that sends v to v if $v \in W'$ and sends v to $v + s$ if $v \in W$. The restrictions of φ both to $I \cap W'$ and to $I \cap W$ are bijections. Also, $\varphi(I \cap W')$, being $I \cap W'$, is disjoint from W , and $\varphi(I \cap W) \subseteq W$, this latter since the vertices of W with a neighbor outside W are all unoccupied. This shows $|I_s| = |I|$. To see that I_s is an independent set, note first that $\varphi(I \setminus W)$ and $\varphi(I \cap W)$ are both independent sets, so we need only rule out the possibility of having $v_1 \in I \cap W$ and $v_2 \in I \cap W'$ with $\varphi(v_1)\varphi(v_2)$ an edge in \mathbb{Z}^2 . Since (as we have already observed) $\varphi(v_1) \in W$, and $\varphi(v_2) \in W'$, such an adjacency would put $\varphi(v_1)\varphi(v_2)$ in γ ; but since $\varphi(v_2) = v_2$ this would lead to an edge in γ with one endvertex occupied, contradicting (5).

Next we show that \tilde{I}_s is disjoint from I_s . It is clearly disjoint from $\varphi(I \cap W')$. It is also easily seen to be disjoint from $\varphi(I \cap W)$, since all vertices v in $\varphi(I \cap W)$ have $v - s \in W$, and no vertices in \tilde{I}_s have this property. Finally we need to show that no $v \in \tilde{I}_s$ is adjacent to something in I'_s . There cannot be a $w \in \varphi(I \cap W')$ with $vw \in \mathbb{Z}^2$, for vw would then be in γ and have one endvertex (w) occupied. Next we consider a $w \in \varphi(I \cap W)$ with vw an edge of \mathbb{Z}^2 . We cannot have $w = v + s$, for then we would have $v \in I$ (again creating an edge in γ with one endvertex, this time v , occupied). We cannot have $w = v - s$ since this would put w into W' (by definition of \tilde{I}_s), and we know $w \in W$ since $\varphi(I \cap W) \subseteq W$. There remains the case $w = v + s'$, with s' perpendicular to s . But in this case, $w - s$ is an occupied vertex in W , and $v - s$ is a neighbor of $w - s$ that is outside W , again creating an impossible edge in γ . \square

Proof of Lemma 3.6: Γ determines I and so W , and this together with s determines \tilde{I}_s (which, crucially, depends only on W and s and not on I). This allows S to be determined,

as $S = J \cap \tilde{I}_s$, from which I_s can be determined as $I_s = J \setminus S$. Finally we determine I as $I = (I_s \cap W') \cup \{\sigma_s^{-1}(v) : v \in I_s \cap W\}$. \square

Proof of Lemma 3.7: Because I is m -odd we know that no even vertex of U_m is in I , and nor is any even vertex outside U_m that is adjacent to something in U_m . We aim to establish that after the shift operation no odd vertex of U_m is in I , and nor is any odd vertex outside U_m that is adjacent to something in U_m ; this shows that I_s is m -even, and since in going from I_s to I_s'' we only add even vertices, so also is I_s'' .

That no odd vertex of U_m is in I after the shift is clear, since R includes $\sigma_{-s}(v)$ for every odd $v \in U_m$, no such $\sigma_{-s}(v)$ is in I , and the status of v with regards membership of I_s is identical to the status of $\sigma_{-s}(v)$ with regards membership of I . The same argument holds for any odd vertex v outside U_m adjacent to something in U_m for which either $\sigma_{-s}(v) \in U_m$ or $\sigma_{-s}(v)$ is outside U_m but adjacent to something in U_m .

There remains the case of odd v , adjacent to something in U_m , with $\sigma_{-s}(v)$ not in U_m and not adjacent to something in U_m . If such a v is not in I , then it is clearly not in I_s . If $v \in I$ then $v \in R$ and $\sigma_{-s}(v) \notin I$ and so as before $v \notin I_s$. \square

3.4 Contours as taxi walks

In this section we establish the connection between contours and taxi walks, which allows us to give the proof of Lemma 3.4. The key ingredient is the following.

Lemma 3.9

1. Viewed as a polygon in \mathbb{Z}_\diamond^2 , if Γ turns, goes straight for an odd number of steps, and turns again, then the second turn must be in the same direction as the first, while if it goes straight for an even number of steps, then the second turn must be in the opposite direction.
2. If Γ turns, then it cannot turn again after a single step.

Proof: We first show that if Γ turns, goes straight for an odd number of steps, and turns again, then the second turn must be in the same direction as the first, while if it goes straight for an even number of steps, then the second turn must be in the opposite direction.

Suppose that $(e_1, f_1, f_2, \dots, f_{2k+1}, e_2)$ is a list of consecutive edges in Γ , with e_1 perpendicular to f_1 , all the f_i 's parallel, and e_2 perpendicular to f_{2k+1} . Without loss of generality e_1 is the edge in Γ from $(-1/2, 1)$ to $(0, 1/2)$, which by Lemma 3.8 forces f_1 to go from $(0, 1/2)$ to $(1/2, 1)$, which forces f_i to go from $(0, 1/2) + (i-1)(1/2, 1/2)$ to $(0, 1/2) + i(1/2, 1/2)$, and in particular f_{2k+1} to go from $(k, k + 1/2)$ to $(k + 1/2, k + 1)$. Again by Lemma 3.8, e_2 must now go from $(k + 1/2, k + 1)$ to $(k, k + 3/2)$. This shows that two turns in Γ separated by an odd number of steps must both go in the same direction (counterclockwise in this case). The case of Γ taking an even number of steps between turns is dealt with similarly.

Next we consider the possibility of Γ taking two consecutive turns. Suppose that the turns are taken around vertex v , in the sense that v has neighbors (read off in cyclic order) a, b, c, d , and Γ has edges from av to bv , from bv to cv , and from cv to dv . (Bearing Lemma 3.8 in mind, no other situation is possible).

Consider first the case where v is an odd vertex. In the 1-by-1 square that a , v and d are three corners of, the construction of Γ dictates that there must be an edge of Γ from av to dv (we are either in case iii or case iv). This means that Γ encloses just the odd vertex v , which, by Lemma 3.1, cannot happen.

If v is even, let q be the vertex that completes the 1-by-1 square that includes b and c as corners, r the one for c and d , s the one for a and b , and t the one for a and d . Note that Γ cannot have an edge from av to dv , for if it did it would be a 4-cycle enclosing a single vertex v , implying that $|R| = 1$, a contradiction since $U_m \subseteq R$. Looking at the construction rules for Γ , we see that we must have $q \notin R$, $r \notin R$, $s \notin R$ and $t \in R$, and Γ has edges from av to at and from dv to dt . Note now that $v \in R$ has all four of its neighbors outside R . This is not possible, since by construction of R every even vertex of R must have a neighbor in R . It follows that Γ cannot take two consecutive turns. \square

The import of Lemma 3.9 is that Γ can be thought of as a taxi walk; we now record this key fact formally.

Lemma 3.10 *Let m_x, m_y be integers such that $(m_x, m_y - 1/2)$ is the apex of a “vee” in Γ ; that is to say, $(m_x - 1/2, m_y)$ is adjacent to $(m_x, m_y - 1/2)$ in Γ , and $(m_x, m_y - 1/2)$ is adjacent to $(m_x + 1/2, m_y)$. There is a unique orientation of the edges of \mathbb{Z}_\diamond^2 such that it becomes isomorphic to $\vec{\mathbb{Z}}^2$ via a translation that sends $(m_x, m_y - 1/2)$ to the origin, followed by a clockwise rotation through $\pi/4$, followed by a dilation by $\sqrt{2}$. Under this orientation, if the edge from $(m_x - 1/2, m_y)$ to $(m_x, m_y - 1/2)$ is removed from Γ then the residue is mapped to a taxi walk of length $|\Gamma| - 1$.*

Proof of Lemma 3.4: By Lemma 3.10 an element Γ of \mathcal{C}_ℓ^m is fully described by specifying a midpoint $(m_x, m_y - 1/2)$ (m_x, m_y integers) of an edge in \mathbb{Z}^2 where Γ makes a “vee” turn, followed by specifying a taxi walk of length $4\ell - 1$. Since Γ is a simple closed curve of length $4\ell/\sqrt{2}$ that encloses the origin, there are at most $O(\ell^2)$ choices for the pair (m_x, m_y) , and by Lemma 2.2 there are at most $f_{\text{taxi}}(4\ell - 1)\mu_{\text{taxi}}^{4\ell-1}$ choices for the taxi walk. The lemma follows. \square

4 The taxi walk connective constant (Theorem 2.2)

In this section we prove Theorem 2.2. A helpful initial observation is that $\vec{\mathbb{Z}}^2$ is vertex-transitive; specifically, for each $(x, y) \in \mathbb{Z}^2$, the bijective map $f_{(x,y)} : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$ given by

$$f_{(x,y)} = \begin{cases} \text{translation by } (-x, -y) & \text{if } x, y \text{ both even} \\ \text{translation by } (-x, -y), \text{ then rotation through } \pi \text{ radian} & \text{if } x, y \text{ both odd} \\ \text{translation by } (-x, -y), \text{ then reflection across } x\text{-axis} & \text{if } x \text{ odd, } y \text{ even} \\ \text{translation by } (-x, -y), \text{ then reflection across } y\text{-axis} & \text{if } x \text{ even, } y \text{ odd} \end{cases}$$

induces an orientation-preserving bijection of $\vec{\mathbb{Z}}^2$ that sends (x, y) to the origin.

We begin the proof of Theorem 2.2 by establishing the submultiplicativity of c_n (or, equivalently, the subadditivity of $\log c_n$).

Lemma 4.1 *For $n, m \geq 1$, $c_{n+m} \leq c_n c_m$.*

Proof: If we split a taxi walk of length $n + m$ into two pieces, an initial segment of length n and a terminal segment of length m , then both resulting pieces are self-avoiding. Moreover the initial segment of length n is a taxi walk of length n , while the terminal segment of length m gets mapped to a taxi walk of length m by the map $f_{(x,y)}$ described above, where (x, y) is the initial vertex of the terminal segment. It is straightforward to verify that this gives rise to an injective mapping from taxi walks of length $n + m$ to ordered pairs of taxi walks, the first of length n and the second of length m , so that $c_{n+m} \leq c_n c_m$. \square

It follows from Lemma 4.1 that $d_n := \log c_n$ is subadditive, i.e., $d_{n+m} \leq d_n + d_m$. By Fekete's Lemma (see, e.g., [27, Lemma 1.2.2]) we know that $\lim_{n \rightarrow \infty} d_n/n$ exists and that

$$\lim_{n \rightarrow \infty} \frac{d_n}{n} = \inf_n \frac{d_n}{n}. \quad (10)$$

Thus we can write the number of taxi walks of length n as $c_n = f_{\text{taxi}}(n) \mu_{\text{taxi}}^n$, where μ_{taxi} is a constant and $f_{\text{taxi}}(n)$ is subexponential in n .

We have already (in Section 2) observed that $\sqrt{2} \leq \mu_{\text{taxi}} \leq (1 + \sqrt{5})/2$. Various techniques from the self-avoiding walk literature — subadditivity, Alm's method, the Goulden-Jackson cluster method, and Kesten's methods of bridges and irreducible bridges — can be used to improve both bounds. We now discuss these methods and our associated results.

4.1 Upper bounds on μ_{taxi}

Subadditivity gives us a strategy for getting a better upper bound on μ_{taxi} . From (10) we see that for all n , $\log c_n/n$ is an upper bound for $\log \mu_{\text{taxi}}$. Then, using that $c_{60} = 2189670407434$ (see [6]) gives the bound $\mu_{\text{taxi}} < 1.60574$ and $\mu_{\text{taxi}}^4 - 1 < 5.6482$.

The connective constant for ordinary self-avoiding walks has been well studied, and some of the methods used to obtain bounds there can be adapted to deal with taxi walks. In this section, we adapt two methods due to Alm [1] and Goulden and Jackson [16] to bound μ_{taxi} and thus establish Theorem 2.2. The bounds derived using these methods are very similar, so both are provided for completeness.

First we discussed the method of Alm [1]. Fix $n > m > 0$. Construct a square matrix $A(m, n)$ whose ij entry counts the number of taxi walks of length n that begin with the i th taxi walk of length m , and end with the j th taxi walk of length m , for some fixed ordering of the walks of length m (formally we mean that if (x, y) is the vertex that begins the terminal segment of length m of the walk, then the map $f_{(x,y)}$ described earlier sends the the terminal segment of length m to the j th taxi walk of length m). Then a result of Alm [1, Theorem 1] says that

$$\mu_{\text{taxi}} \leq \lambda_1(A(m, n))^{1/(n-m)},$$

where λ_1 indicates the largest positive eigenvalue. (Note that when $m = 0$ this recovers the subadditivity bound discussed earlier). Alm's result as stated in [1] only applies to bound the ordinary connective constant of a finitely generated lattice, directed or otherwise. His proof is easily seen to go through without change, however, when the extra condition is added that walks do not take two consecutive turns. We calculated $A(20, 60)$; this is a square matrix of dimension 20114 for which we can estimate its largest eigenvalue using MATLAB. This gives that $\mu_{\text{taxi}} < 1.58834$ and $\mu_{\text{taxi}}^4 - 1 < 5.3646$ (again see [6] for this data).

Our second approach is the Goulden-Jackson cluster method [16]. This is an algorithm which takes as input a finite alphabet A , an integer n and a finite list \mathcal{M} of words over A — the elements of which we refer to as *mistakes* — and outputs the number ℓ_n of words of length n over A that do not contain any mistakes as subwords (that is, as strings of consecutive letters in the word).

Recall from the proof of Lemma 2.1 that a taxi walk of length n may be encoded by a pair (a, σ) , where $a \in \{N, E\}$ and σ is a word of length $n-1$ over alphabet $\{s, t\}$. Suppose that \mathcal{M} is a finite set of subwords that is not allowed to occur in any word σ over alphabet $\{s, t\}$ that occurs in an encoding of a taxi walk (for example, tt is one such subword). If ℓ_n is as defined in the last paragraph then we have $c_{n+1} \leq 2\ell_n$, so by subadditivity $\mu_{\text{taxi}} \leq (2\ell_n)^{1/(n+1)}$.

We can improve this slightly. Alm [1, Remark 9] observes that in a vertex-transitive lattice for which any self-avoiding walk of length 1 can be mapped on to any other self-avoiding walk of length 1 by some orientation-preserving symmetry (built from translations, rotations and reflections), the connective constant is bounded above by $(f(n+1)/f(1))^{1/n}$, where $f(m)$ is the number of self-avoiding walks of length m starting for some fixed vertex (by vertex-transitivity, it does not matter which). Applying this to the present situation (where Alm's condition is certainly satisfied, with $f(1) = 2$), we get

$$\mu_{\text{taxi}} \leq \ell_n^{1/n}. \quad (11)$$

If (a, σ') encodes a walk in the Manhattan lattice that takes no two consecutive turns, starts and ends at the origin, and otherwise does not visit any vertex twice, then it is evident that the word σ' cannot occur as a subword of σ in any taxi walk (a, σ) . We refer to such a σ' as a *taxi polygon* of length $|\sigma'| + 1$, where $|\sigma'|$ is the number of letters in σ' . For example, $sstsstsstss$ is a taxi polygon of length 12, and $tstsstsssstsssstst$ is a taxi polygon of length 20.

We have enumerated taxi polygons of length at most 48 (there are 8,009,144 of them). We then used an implementation of the Goulden-Jackson cluster method due to Noonan and Zeilberger [19] to calculate a_{802} with \mathcal{M} consisting of the set of all taxi polygons of length at most 44 (there are 1,721,326 of them) together with the word tt . Via (11) this leads to $\mu_{\text{taxi}} < 1.58746$ and $\mu_{\text{taxi}}^4 - 1 < 5.3506$, as stated in Theorem 2.2. (See [6] for the data and the computer code used to generate it.)

4.2 Lower bounds on μ_{taxi}

To improve the trivial lower bound $\sqrt{2} \leq \mu_{\text{taxi}}$ we consider *bridges* (introduced for ordinary self-avoiding walks by Kesten [17]). A *bridge*, for our purposes, is a taxi walk that begins by moving from the origin $(0, 0)$ to the vertex $(1, 0)$, never revisits the y -axis, and ends by taking a step parallel to the x -axis to a vertex on the walk that has maximum x -coordinate over all vertices in the walk (but note that this maximum does not have to be uniquely achieved at the final vertex).

Let b_n be the number of bridges of length n (by convention $b_0 = 1$). Observe that bridges are supermultiplicative, that is, $b_{n+m} \geq b_n b_m$ (and $\log b_n$ is superadditive). To see this, consider bridges β_1 of length n and β_2 of length m . By the definition of a bridge, it is straightforward to verify that if we concatenate β_1 and the image of β_2 under the map $f_{(x,y)}^{-1}$,

where (x, y) is the terminal vertex of β_1 , then the result is a bridge. Moreover, the map just described from pairs of bridges, the first of length n and the second of length m , to bridges of length $n + m$, is injective. It follows that there are at least b_n^k taxi walks of length kn (just concatenate k length n bridges), so that

$$\mu_{\text{taxi}} = \lim_{m \rightarrow \infty} c_m^{1/m} \geq \lim_{k \rightarrow \infty} (b_n^k)^{1/nk} = b_n^{1/n}.$$

Since $b_{60} = 80312795498$ (see [6]), we get that $\mu_{\text{taxi}} > 1.51965$ and $\mu_{\text{taxi}}^4 - 1 > 4.3330$.

Using Kesten's more sophisticated notion of *irreducible bridges* (bridges that are not the concatenation of shorter bridges), we can get significantly better bounds. Our discussion follows the approach of Alm and Parviainen [2]).

Say that an internal vertex (x, y) along a bridge is a *cutvertex* if the walk from the origin up to (x, y) is a bridge, after (x, y) the next vertex of the walk is $(x + 1, y)$, and the walk from (x, y) to the end (more correctly, the image of this walk under $f_{(x, y)}$) is also a bridge. Say that a bridge is *irreducible* if it does not have a cutvertex. Denote by a_n the number of irreducible bridges of length n (by convention $a_0 = 0$).

Fix $n \geq 1$. For each $\ell \geq 1$, each solution to $k_1 + \dots + k_\ell = n$ with each $k_i \geq 1$, and each sequence of irreducible bridges (p_1, \dots, p_ℓ) with p_i of length k_i for each i , there corresponds a bridge of length n obtained by concatenating the p_i 's (with suitable translations, reflections and rotations where necessary). Moreover, each bridge of length n is obtained exactly once in this process. It follows that

$$b_n = \sum_{\ell \geq 1} \sum \left\{ \prod_{i=1}^{\ell} a_{k_i} : \text{compositions } k_1 + \dots + k_\ell = n \right\},$$

and so, setting $B(x) = \sum_{k \geq 0} b_k x^k$ and $A(x) = \sum_{\ell \geq 1} a_\ell x^\ell$, we have

$$B(x) = \frac{1}{1 - A(x)}. \quad (12)$$

Notice that $A(x) = 1$ has a unique solution r_{pos} in the interval $(0, 1)$. Let r be any upper bound on r_{pos} . From (12), standard facts about generating functions (see e.g. [30, Section 2.4]) tells us that r_{pos} (and thus r) is an upper bound on the radius of convergence of $B(x)$; consequently

$$\limsup_{n \rightarrow \infty} b_n^{1/n} \geq 1/r.$$

But we also know that $c_n \geq b_n$, implying that

$$\mu_{\text{taxi}} = \limsup_{n \rightarrow \infty} c_n^{1/n} \geq \limsup_{n \rightarrow \infty} b_n^{1/n} \geq 1/r.$$

It follows that an upper bound on r_{pos} , the unique positive solution to $A(x) = 1$, yields a lower bound on μ_{taxi} .

Consider a sequence $(a'_n)_{n=1}^\infty$ with $0 \leq a'_n \leq a_n$ for each n . Set $A'(x) = \sum_{n=1}^\infty a'_n x^n$. As is the case with $A(x)$, the equation $A'(x) = 1$ has a unique solution r' in the interval $(0, 1)$, which moreover clearly satisfies $r' \geq r_{\text{pos}}$. We record the conclusion to this discussion as a theorem.

Theorem 4.2 *With the notation as above, if $x > 0$ satisfies $\sum_{n=1}^{\infty} a'_n x^n > 1$ then $\mu_{\text{taxi}} > 1/x$.*

For any $N \geq 1$, the coefficients of the power series of $1/B(x)$ up to the coefficient of x^N are determined by the coefficients of the power series of $B(x)$ up to x^N , and so using (12) the coefficients of $A(x)$ up to x^N are determined by the coefficients of $B(x)$ up to x^N . We know the coefficients of $B(x)$ up to x^{60} and from this we can easily calculate a_n for $n \leq 60$ (see [6]). Taking $a'_n = a_n$ for $n \leq 60$ and $a'_n = 0$ for $n > 60$ yields $\mu_{\text{taxi}} > 1.55701$ and $\mu_{\text{taxi}}^4 - 1 > 4.8771$, as stated in Theorem 2.2.

5 Concluding remarks

- In an early version of this work [5] we employed the Peierls argument described in Section 3 to separate μ_n^e and μ_n^o by establishing

$$\mu_n^e(E) < 1/3$$

in place of (1), where the event E was defined in terms of fault lines and crosses, as defined in [24]. The idea of replacing E with \mathcal{O}_m , and conditioning on \mathcal{H}_m , which significantly streamlines the analysis, was suggested to us by Kotecký [18] after the third author spoke on this work during the 2013-14 Warwick EPSRC Symposium on Statistical Mechanics. We are very grateful to him for this suggestion.

- The standard Peierls argument for establishing phase coexistence tries to separate μ_n^e and μ_n^o using the event $\{v \in I\}$ where $v \in \mathcal{O}$ is some fixed vertex. Indeed, in [7] it is shown using the FKG inequality that there is phase coexistence for the hard-core model on \mathbb{Z}^2 if and only if $\limsup_{n \rightarrow \infty} \mu_n^e(\{v \in I\}) < \limsup_{n \rightarrow \infty} \mu_n^o(\{v \in I\})$. However, analyzing $\mu_n^e(\{v \in I\})$ using the approach described in the present paper requires considering contours of length 4ℓ for all $\ell \geq 3$. This in turn necessitates controlling the sub-exponential term in the growth rate of taxi walks, which in turn leads to poorer bounds on λ . Using \mathcal{O}_m as the distinguishing event, which ensures that all contours are long, obviates this necessity. This specifically comes in to play with the lower bound on ℓ in (9), coming from Lemma 3.3.
- In [5] it was shown that if λ satisfies $\lambda > \mu_{\text{taxi}}^4 - 1$ then Glauber dynamics for independent sets on the $n \times n$ torus takes time at least e^{cn} to mix, for some constant $c = c(\lambda) > 0$, and that if also $2(1 + \lambda) > \mu_{\text{taxi}}^2(1 + \sqrt{1 + 4\lambda})$ then the same is true for the $n \times n$ grid. Based on the computations in that reference, it was possible to conclude slow mixing for $\lambda > 5.3646$ on the torus, and $\lambda > 7.1031$ on the grid. Using our improved bounds on μ_{taxi} here, we can improve these bounds to $\lambda > 5.3506$ on the torus, and $\lambda > 7.0852$ on the grid.
- One way to improve our lower bound on μ_{taxi} would be to construct families of irreducible bridges of various lengths $n > 60$, and to use the sizes of these families as the a'_n s in Theorem 4.2. So far we have only had slight success with this approach, obtaining $\mu_{\text{taxi}} > 1.55711$ and $\mu_{\text{taxi}}^4 - 1 > 4.8786$. The details are messy, and we choose not to include them here, but they can be found at [6].

References

- [1] S. Alm. Upper bounds for the connective constant of self-avoiding walks. *Combinatorics, Probability & Computing* **2**:115–136, 1993.
- [2] S. Alm and R. Parviainen. Bounds for the connective constant of the hexagonal lattice. *J. Phys. A: Math. Gen.* **37**:549–560, 2004.
- [3] R. J. Baxter, I. G. Enting and S. K. Tsang. Hard-square lattice gas. *Journal of Statistical Physics* **22**:465–489, 1980.
- [4] V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random cluster model is critical above $q = 1$. *Probability Theory and Related Fields* **153**:511–542, 2012.
- [5] A. Blanca, D. Galvin, D. Randall and P. Tetali. Phase Coexistence and Slow Mixing for the Hard-Core Model on \mathbb{Z}^2 . *Lecture Notes in Comput. Sci.* **8096**:379–394, 2013 (*Proc. APPROX/RANDOM 2013*).
- [6] A. Blanca, E. Chen, D. Galvin, D. Randall and P. Tetali. Taxi walks computations. <http://nd.edu/~dgalvin1/TD/>
- [7] J. van den Berg and J.E. Steif. Percolation and the hard-core lattice model. *Stochastic Processes and their Applications* **49**: 179–197, 1994.
- [8] C. Borgs. Personal communication.
- [9] C. Borgs, J.T. Chayes, A. Frieze, J.H. Kim, P. Tetali, E. Vigoda, and V.H. Vu. Torpid mixing of some MCMC algorithms in statistical physics. *Proc. 40th IEEE Symp. on Foundations of Computer Science (FOCS)* 218–229, 1999.
- [10] G. Brightwell and P. Winkler. Graph homomorphisms and phase transitions. *J. Combin. Theory Ser. B* **77**: 221–262, 1999.
- [11] R.L. Dobrushin. The problem of uniqueness of a Gibbs random field and the problem of phase transitions. *Functional Analysis and its Applic.* **2**: 302–312, 1968.
- [12] R. Dobrushin. The description of a random field by means of conditional probabilities and conditions of its regularity. *Theory Probab. Appl.* **13**: 197–224, 1968.
- [13] D. Galvin. Sampling independent sets on the discrete torus. *Random Str. & Algorithms* **33**: 356–376, 2008.
- [14] D. Galvin and J. Kahn. On phase transitions in the hard-core model on \mathbb{Z}^d . *Combinatorics, Probability & Computing* **13**: 137–164, 2004.
- [15] H.-O. Georgii. Gibbs Measures and Phase Transitions. de Gruyter, Berlin, 1988.
- [16] I. Goulden and D.M. Jackson. An inversion theorem for cluster decompositions of sequences with distinguished subsequences. *J. London Math. Soc.* **20**: 567–576, 1979.

- [17] H. Kesten. On the number of self-avoiding walks. *J. Math. Phys.*, **4**: 960–969, 1963.
- [18] R. Kotecký. Personal communication.
- [19] J. Noonan, D. Zeilberger. The Goulden-Jackson Cluster Method: Extensions, applications, and implementations. *J. Difference Equations Appl.*, **5**:355–377, 1999.
- [20] L. Onsager. Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Physics Review Letters* **65**: 117–149, 1944.
- [21] J. Pantone. Personal communication.
- [22] J. Pantone. The method of differential approximants for enumerative combinatorics. Forthcoming.
- [23] D. C. Radulescu and D. F. Styer. The Dobrushin-Shlosman phase uniqueness criterion and applications to hard squares. *Journal of Statistical Physics* **49**: 281–295, 1987.
- [24] D. Randall. Slow mixing of Glauber dynamics via topological obstructions. *Proc 17th ACM-SIAM Symp. on Discrete Algorithms (SODA)* 870–879, 2006.
- [25] R. Restrepo, J. Shin, P. Tetali, E. Vigoda and L. Yang. Improved mixing condition on the grid for counting and sampling independent sets. *Probab. Theory Relat. Fields* **156**:75–99, 2013.
- [26] A. Sinclair, P. Srivastava, D. Štefankovič and Y. Yin. Spatial mixing and the connective constant: Optimal bounds. *Probab. Theory Relat. Fields*, 2016. doi:10.1007/s00440-016-0708-2.
- [27] J.M. Steele. Probability Theory and Combinatorial Optimization. SIAM, 1997.
- [28] J. C. Vera, E. Vigoda and L. Yang. Improved Bounds on the Phase Transition for the Hard-Core Model in 2-Dimensions. *Lecture Notes in Computer Science* **8096**: 699–713, 2013.
- [29] D. Weitz. Counting independent sets up to the tree threshold. *Proc. 38th ACM Symp. on the Theory of Computing (STOC)* 140–149, 2006.
- [30] H. Wilf. generatingfunctionology. A.K. Peters, 2005.